Error Correction, Sensory Prediction, and Adaptation in Motor Control

The Harvard community has made this article openly available. **Please share** how this access benefits you. Your story matters

Citation	Shadmehr, Reza, Maurice A. Smith, and John W. Krakauer. Forthcoming. Error correction, sensory prediction, and adaptation in motor control. Annual Review of Neuroscience 33.
Published Version	http://arjournals.annualreviews.org/doi/abs/10.1146/annurev-neuro-060909-153135
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:3874487
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

ERROR CORRECTION, SENSORY PREDICTION, AND ADAPTATION IN MOTOR CONTROL

Reza Shadmehr[1], Maurice A. Smith[2], and John W. Krakauer[3]

[1]Dept. of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, MD 21205

[2]School of Engineering and Applied Science, Harvard University, Cambridge, MA 02138

[3]The Neurological Institute, Columbia University School College of Physicians and Surgeons, New York, NY 10032

To appear in: *Annual Reviews of Neuroscience*

Key words

Forward models, reaching, saccades, motor adaptation, error feedback, sensorimotor integration.

Contact: Reza Shadmehr, Johns Hopkins School of Medicine, 720 Rutland Ave, Baltimore, MD 21205. Email: shadmehr@jhu.edu phone: 410-614-2458.

Disclosure statement

The authors are not aware of any affiliations, membership, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

Acknowledgements

ABSTRACT

Motor control is the study of how organisms make accurate goal-directed movements. There are two problems that the motor system must solve in order to achieve such control. The first problem is that sensory feedback is noisy and delayed, which can make movements inaccurate and unstable. The second problem is that the relationship between a motor command and the movement it produces is variable, as the body and the environment can both change. A solution is to build adaptive internal models of the body and the world. The predictions of these internal models, called forward models because they transform motor commands into sensory consequences, can be used to both produce a lifetime of calibrated movements, and to improve the ability of the sensory system to estimate the state of the body and the world around it. Forward models are only useful if they produce unbiased predictions. Evidence shows that forward models remain calibrated through motor adaptation: learning driven by sensory prediction errors.

We have the ability to control our movements during a vast array of behaviors ranging from making simple limb movements to dribbling a basketball, throwing a baseball, and juggling. We marvel at the accomplishments of athletes, yet, from a theoretical standpoint, even the ability to make the simplest eye and arm movements accurately is quite extraordinary. Consider that the motors that actuate robots reliably produce the same force for a given input. Yet, our muscles quickly fatigue, altering their responses from one movement to the next. The sensors that record motion of a robot do so with far more precision than one finds in the response of our proprioceptive neurons. The transmission lines that connect a robot's motors and sensors to controller move information at the speed of light, and the controller can process sensory information to issue commands in microseconds. In contrast, our transmission lines (axons) move information slower than the speed of sound, and neural computations often require tens of milliseconds. Therefore, our ability to produce a lifetime of accurate movements is not because we are born with an invariant set of actuators, precise set of sensors, or fast transmission lines, but rather because we are born with a nervous system that adapts to these limitations and continuously compensates for them. If left uncompensated, these inherent limitations could give rise to systematic errors in our movements. How the brain is able to predict and correct systematic errors to produce a lifetime of accurate movements is the subject of this review.

Correcting movement errors without sensory feedback

A typical saccade takes less than 80ms to complete and moves the eyes at more than 400 deg/s. Such movements are too brief for visual or proprioceptive feedback to influence control of the eyes during the saccade (Keller and Robinson, 1971;Guthrie et al., 1983). However, a fundamental problem is that the motor commands that initiate the movement are highly variable and this variability is related to the context in which the eye movement is made. For example, people make saccades with higher velocities in anticipation of seeing a more interesting visual stimulus (e.g., image of a face) (Xu-Wilson et al., 2009b). If a target is presented and they are instructed to look in the opposite direction, saccade velocities are much lower than if they are asked to look to the target (Smit et al., 1987). Repeating a visual target (Straube et al., 1997;Chen-Harris et al., 2008;Golla et al., 2008) or reducing the reward associated with that stimulus (Takikawa et al., 2002) also reduces saccade velocities. On the other hand, increasing the reward associated with the target (Takikawa et al., 2002), making the target the goal of both the eye and the arm movements (van Donkelaar, 1997;Snyder et al., 2002), or unexpectedly

changing the characteristics of the target (Xu-Wilson et al., 2009a) all result in increased saccade velocities without altering saccade amplitude. Thus, the motor commands that accelerate the eyes toward a given target are affected by the expected reward, attention, or cognitive state of the subject. These factors all induce variability in the motor commands that start the movement. Despite this variability, the brain accurately guides the eyes to the target without sensory feedback.

Another problem is that we sometimes blink when we make saccades. As the eyelid comes down on the eyes, it acts as a source of mechanical perturbation that pushes the eye, altering its trajectory. Remarkably, the motor commands that guide the eyes during a blink-affected saccade appear to take into account this self-induced source of perturbation: the eyes tend to arrive on target (Rottach et al., 1998). How does the brain take into account these potential sources of variability without the use of sensory feedback?

Some three decades ago David Robinson proposed that endpoint accuracy of saccades is possible because the brain incorporates an internal feedback process through the cerebellum that monitors the motor commands and corrects them online (Robinson, 1975). Corollary discharge encodes a copy of the motor command (so-called efference copy), and this efference copy could be processed to predict the consequences of actions before sensory feedback is available. Predicting the consequences of a motor command is called a *forward model*. In the current computational view of motor control (e.g., Shadmehr and Krakauer (2008)), the cerebellum may be a forward model (Pasalar et al., 2006) that uses efferent copy to predict consequences of motor commands and to correct the movement as it is being generated (Fig. 1A). For example, if the internal feedback is intact, variability in the commands that initiate the saccade might be compensated via cerebellar dependent commands that arrive later during the same saccade.

A strong prediction of this idea is that if an experiment could impose variability in the commands that accelerate the eyes, a subject who is missing this internal feedback process (e.g., cerebellar patients) would exhibit an inability to compensate for that variability, resulting in dysmetric saccades. A simple way to induce variability in the motor commands that accelerate the eyes is to present the visual targets in a repeating pattern: in healthy subjects, this repetition results in a decline in saccade velocities without affecting saccade amplitudes (Fuchs and Binder, 1983;Straube et al., 1997). The origin of this decline is not well understood. However, the decline is not due to muscle fatigue but likely of neural origin (Xu-Wilson et al., 2009a). One

possibility is that the repetition of the stimulus acts to de-value the target of the movement, and a reduced value associated with the target results in reduced saccade velocities (Takikawa et al., 2002). Golla et al. (2008) and Xu-Wilson et al. (Xu-Wilson et al., 2009a) used this repetition method to introduce variability in the motor commands that initiated saccades. They then examined the ability of patients with cerebellar damage to respond to this variability. Both groups of investigators found that repetition of the stimulus produced strong reductions in saccade velocities in the healthy and cerebellar subjects (Fig. 1B). That is, if saccade velocities are viewed as a proxy for the value that the brain assigns the target (Takikawa et al., 2002;Xu-Wilson et al., 2009b), then in both healthy and cerebellar subjects the repetition of the target resulted in its devaluation. However, whereas in the healthy subjects saccade amplitudes remained accurate, in the cerebellar patients the saccades fell short of the target as the velocities declined (Fig. 1B). It appeared that in healthy people, the variability in the motor commands that initiated the saccade was generally compensated via motor commands that arrived later in the same saccade. However, the compensation was missing in cerebellar subjects.

In summary, there is variability in the motor commands that initiate even the simplest movements like saccades. If left uncompensated, this variability would result in dysmetric movements. The brain appears to monitor the motor commands and compensate for the variability, a process that is consistent with that of an internal model that predicts the sensory consequences of motor commands.

Correcting movement errors with delayed sensory feedback

Unlike saccade, most movements are long enough in duration that sensory feedback plays an essential role in their control. All goal-directed arm movements fall into this category. However, the problem is that the delays inherent in sensory feedback can destabilize movements, confining online error correction to peripheral mechanisms characterized by fast spinal reflexes and intrinsic biomechanical properties of the muscles. Such a control strategy, however, allows for only a very narrow class of error feedback behaviors. Despite these long delays, supra-spinal, long-latency responses to perturbations are often much larger in amplitude than their short-latency counterparts (Strick, 1978), and response to visual and proprioceptive error are salient throughout the entire course of even relatively quick arm movements (Saunders and Knill, 2003;Saunders and Knill, 2004). Together these observations suggest that cortically-modulated, in addition to spinal, error correction mechanisms play a role in even short, rapid arm movements.

This ability to perform cortically-driven online error feedback control is remarkable in light of the long sensorimotor delays we experience.

Why do sensorimotor loop delays present a problem for real-time feedback control? Consider a simple example. Most of us have experienced the trauma of being scalded in an unfamiliar shower. This situation nicely illustrates how delayed feedback can lead to unstable control. In an unfamiliar bathroom, an unsuspecting/naïve individual may start his shower at too low of a temperature and then want to increase the heat. Since the delay between adjusting the temperature control knob and feeling its effects is notoriously long, initial adjustment of the knob has no immediate effect on the water temperature, which may spur the victim-to-be to continue turning up the heat. By the time that the water temperature starts to respond, the heat adjustment may already be at scalding levels. Then when the water becomes too hot, the victim turns down the temperature control. However the temperature continues to rise. The victim responds by turning down the temperature control even farther, and when the temperature finally responds, he is soon freezing. If he were to continue with the same pattern behavior, the victim would continue to experience large, unstable fluctuations in water temperate resulting in repeated freeze / burn cycles. Most such victims soon decide to stop adjusting the temperature control continually - abandoning rapid, real-time feedback adjustments for a prolonged "wait & see" approach after each small temperature adjustment. This avoids large temperature fluctuations at the expense of an extended time for achieving a comfortable water temperature. Note that the whole traumatic scenario could have been avoided if the victim exercised a good understanding of how the water temperature would react to each adjustment of the knob, as would be the case in his own shower. Such an understanding entails the ability to predict the sensory consequences of motor actions, i.e., a forward model.

If the motor system has the means to predict the sensory state of the motor apparatus, instabilities arising from delays in measuring that state can be effectively eliminated (Miall and Wolpert, 1996;Bhushan and Shadmehr, 1999). Long-latency online feedback control, therefore, might rely on a forward model of dynamics, which would enable the motor system to predict state variables such as position and velocity based on a history of motor commands (Mehta and Schaal, 2002;Ariff et al., 2002;Flanagan et al., 2003). If this predicted state is an unbiased estimate of the actual state, then feedback control driven by these predicted states could essentially take place in real-time and avoid the instability normally associated with feedback delays.

Feedback responses driven by internal models can also be more effective in compensating for a perturbation than responses driven solely by sensory feedback alone. For example, when the arm is moved because of an unexpected perturbation, the short latency spinal reflexes respond solely to the muscle stretch, but longer latency reflexes produce a response that also takes into account the consequences of the net torques on the joints (Lacquaniti and Soechting, 1984;Soechting and Lacquaniti, 1988;Kurtzer et al., 2008). This is important because a perturbation that stretches one muscle often cannot be compensated by activating that muscle alone. That is, because of intrinsic mechanical coupling between the physical dynamics of connected joints such as the elbow and the shoulder (so called interaction torques), responding to a stretch of a shoulder muscle by activating that muscle would produce a motion of not only the shoulder, but also of the elbow. Therefore, responses that take into account these interaction torques can counteract the consequences of external perturbations more effectively than responses that are specific to the muscles that were stretched by the perturbation (Kurtzer et al., 2008). These studies suggest that long latency reflexes take into account the physical dynamics of the limb in producing a compensatory response to a perturbation.

One way to test the role of forward models during ongoing arm movements is to have people generate motor commands based on their estimate of current limb position. For example, if the hand is moving and the brain can predict the current location of the hand and its velocity, then the motor commands should reflect this predicted state. If unexpected perturbations to this state are encountered, then the predicted state will not match the actual state. However, after the sensorimotor loop delay, accurate predictions should be restored if the forward model can integrate information about motor commands with the observed sensory feedback. Wagner and Smith (Wagner and Smith, 2008) tested this idea by exposing subjects to novel velocity-dependent dynamics to train a relationship between hand velocity and lateral force. On some trials, the reaching movement was perturbed. In these cases, lateral force profiles, which reflected real-time internal estimates of hand velocity, were initially disrupted but became extremely accurate 150ms after perturbation offset, even though the hand velocity during this period was clearly different than what had been planned. This demonstrates that a forward model is providing accurate state estimates that are then fed to the previously adapted controller to provide corrections to the unexpected perturbation without a time lag.

If the brain had to rely solely on sensory feedback, then the motor commands should reflect a time delayed estimate of hand position rather than a real-time prediction. Miall et al. (2007) used

this idea to test the hypothesis that the cerebellum was involved in predicting the position of the hand during a movement. They had people move their hands laterally until they heard a tone, at which point visual feedback was removed and they reached in a forward direction toward a target location (Fig. 2A). At the time of the tone, the brain needs to generate motor commands that bring the hand from its current position to the target. These commands depend on the estimated state of the hand. If this state estimate is primarily due to a delayed sensory feedback, then the estimate of hand position will be closer to the shoulder than in reality, in which case the motor commands will miss the target to the right. Alternatively, if this state estimate is primarily due to an unbiased prediction, then the estimate of hand position will be near its actual position, in which case the motor commands will bring the hand to the target. When people heard the tone, they brought their hand very near the target (control group, Fig. 2B). However, if the cerebellum could be disrupted by a series of transcranial magnetic stimulation (TMS) pulses right after the tone, the motor commands that would be generated would miss the target to the right. That is, disruption of the cerebellum should make the state estimate 'stale', reflecting not the current position of the hand but a position some 130ms ago. This is indeed what was observed. These results are consistent with the idea that the cerebellum predicts the state of the limb from the history of motor commands, allowing one to act on this estimate of state rather than relying solely on a delayed sensory feedback, and suggests that cerebellar output might provide a 'motion update' signal which can be combined with delayed sensory feedback elsewhere in the brain in order to generate real-time state estimates for motor control.

The effect of predicting sensory consequences of motor commands on perception

A clear advantage of making sensory predictions is that the brain does not have to wait for the sensory measurements before it can act. However, there is a more fundamental advantage to making predictions, and this has to do with perception.

Our ability to estimate the state of our body and the external world appears to be a combination of two streams of information: one in which our brain predicts what should happen, and one in which our sensory system reports what did happen. The advantage of this is that if our predictions are unbiased, then our perception (and the decisions that are made based on that perception) will be better than if we had to rely on sensory measurements alone. In a sense, our perception will be more accurate (e.g., display less variance), if we combine what we predicted with what our sensory system measured.

Although this may seem like a fairly new idea, it was first proposed about a thousand years ago (around 1011) by Ibn al-Haytham, an Iraqi scientist (also known as Alhazen), in his *Book of Optics* in which he explained the 'moon illusion', the perception that the moon is bigger when it is near the horizon than high in the sky, in terms of expectations about distance (although we perceive it this way, when measured by a simple camera the size of the moon is actually a bit smaller near the horizon, as it is farther away by about the radius of the earth). In 1781 Immanuel Kant in his theory of *idealism* claimed that our perceptions are not the result of a physiological process in which, for example, the eyes faithfully transmit visual information to the brain, but rather, our perceptions are a result of a psychological process in which our brain combines what it already believes with the sensory information. He wrote "The understanding can intuit nothing, the senses can think nothing. Only through their union can knowledge arise." (I. Kant, *Critique of Pure Reason*, trans. N. K. Smith, 1781; New York: St. Martin's Press, 1965, p. 93.)

If we follow this line of reasoning and return to our example of predicting the sensory consequences of the motor commands that move the eyes in a saccade, we might guess that during the post-saccadic period, the brain should have a better estimate of eye position than if it simply relied on proprioceptive information (from eye muscles) alone. In theory, the combination of the two streams of information (prediction and sensory feedback) should allow it to estimate eye position better. As a consequence, subsequent movements that depend on this position estimate should be more accurate.

For example, consider a task in which subjects are asked to make two saccades in succession in the dark so that the endpoint of the first saccade is the starting point of the second saccade. The motor commands that generate the second saccade should take into account the consequences of the motor commands for the first saccade. That is, planning of the second saccade can benefit from the efference copy of the first saccade.   Indeed, a neural pathway for the transmission of this efference copy information appears to be from the superior colliculus, through the mediodorsal nucleus in the thalamus, to the frontal eye fields (Sommer and Wurtz, 2002;Sommer and Wurtz, 2006). When two eye movements are made in rapid succession, the motor command for the second saccade takes the first into account, but when this pathway is disrupted, this ability is markedly reduced (Sommer and Wurtz, 2002).

Usually when we make saccades, visual information is present at the end of each saccade. This visual information is simply a *remapped* version of the visual field before the saccade. In effect, the eye is a camera that moves during a saccade, altering the retinal image. Therefore, the change in the visual information is a predictable event, and is part of the sensory consequences of the motor commands that moved the eye. Accordingly, neurons in posterior parietal cortex have been shown to remap their receptive fields based on such predictions before saccades occur (Duhamel et al, 1992). An interesting prediction is that the brain should have a better estimate of what it sees after a saccade than if that same visual information was only provided passively. Indeed, this was recently confirmed in an experiment in which people were asked to reach to targets that they could see but not predict, vs. targets that they could both predict their location (because of self-generated eye movements, i.e., remapping) and see (Vaziri et al., 2006). Being able to predict the sensory consequences of a saccade allowed subjects to have a more accurate estimate of the location of the target.

Therefore, by predicting the sensory consequences of motor commands, the brain can not only over-come delay in sensory feedback, but perhaps more importantly, it can actually sense the world better than is possible form sensory feedback alone. The latter comes about when our brain combines what it has predicted with what it has measured - two sources of information, when used properly, are better than one.

If our brain could not accurately predict sensory consequences of our motor commands, then we would not be able to sense the world around us in a normal way. An example of this is patient RW, a 35 year old man who was described by Haarmeier et al. (1997). RW suffered a stroke in a region covering parts of the parietal and occipital cortex, affecting a part of cortex that receives vestibular input, a location in which cells are sensitive to visual motion. RW complained of vertigo only when his eyes tracked visual objects and not when his eyes were closed. He explained that when he was watching his son run across a field (a condition in which his eyes moved smoothly to follow his son), he would see the boy running, but he would also perceive the rest of the visual scene (e.g., the trees) smoothly moving in the opposite direction. Haarmeier et al. (1997) demonstrated that when RW moved his eyes, his brain was unable to predict the sensory consequences of the oculomotor commands. As his eyes moved to follow his son, the trees moved in the opposite direction on his retina. The healthy brain predicts that moving the eyes will have the sensory consequence of shifting the image of the stationary world on the retina. We do not perceive this shifting image as real motion of the world because we predict it to be a

consequence of motion of our eyes. In RW, the vertigo was a symptom of the brain's inability to predict such sensory consequences.

You do not need to have a brain lesion to get a feel for what RW sees when he moves his eyes. Take a camera and aim it at a runner and try to move (i.e., 'pan') so that the image of the runner stays at the center of the picture. As you are moving the camera, take a picture. That picture will show a sharply focused runner but a blurry background that appears to be moving in the opposite direction. However, when you are looking at the runner with your naked eyes, the background appears perfectly still. The reason is because your brain predicts the background image shift that should take place on the retina as you move your eyes, and combines it with the actual shift. By combining the observed and predicted images, the parts that agree must have been stationary, and parts that disagree must have moved.

Combining predictions with sensory observations to produce a motor response

To combine two streams of information, one needs to apply a weighting to each stream. In principle, the weight should be higher for the more reliable information source. For example, during a reaching movement, one can make predictions about the state of the arm from the history of motor commands, but as the movement proceeds, one faces the problem of how to combine this prediction with the sensory feedback. Some four decades ago, Rudolph Kalman proposed a principled way that this kind of problem should be solved: combine the two sources of information in a way that minimizes the variance of the resulting estimate (Kalman, 1960). If we view the prediction as the prior belief, and the sensory feedback as the current evidence, then Kalman's algorithm is equivalent to a Bayesian process of integration in which the weight associated with each piece of information depends on the uncertainty of each quantity. Therefore, the idea that emerges is that the brain should not only make predictions about sensory consequences of motor commands, but also incorporate a measure of uncertainty about that prediction.

Körding and Wolpert (Kording and Wolpert, 2004) tested the idea that as a movement took place, the brain combined its predictions about sensory feedback with actual sensory feedback using weights that depended on the uncertainty of each kind of information. They first trained subjects to reach to a goal location by providing them feedback via a cursor on a screen (the hand was never visible). As the finger moved from the start position, the cursor disappeared. Halfway to

the target, the cursor re-appeared briefly. However, its position was, on average, 1 cm to the right of the actual finger position, but on any given trial the actual displacement was chosen from a Gaussian distribution. Because the location of the cursor was probabilistic, the variance of the Gaussian distribution described the confidence with which people could predict the sensory consequences of their motor commands. To control the confidence that the brain should have regarding sensory measurements, they added noise to the display of the cursor: the cursor was displayed as a cloud of dots. On some trials the cursor was shown clearly so the uncertainty regarding its position was low. In other trials the uncertainty was high as the cursor was hidden in a cloud of noise. The idea was that on a given trial, when a subject observes the cursor position midway to the target, she should issue a feedback correction to her movement based on two sources of information: the observation on that trial, and prior prediction regarding where the cursor would have been expected to be. The weighting of each source of error information should be inversely related to the variance of each distribution. Indeed, Körding and Wolpert's (2004) experimental data was consistent with this theoretical framework. Bayesian integration also explains feedback responses to force perturbations (Kording et al., 2004), feedback responses to perturbations to the properties of the target (Izawa and Shadmehr, 2008), and perceptual estimates of position and velocity in the presence of noise (Ernst and Banks, 2002; Weiss et al., 2002; Stocker and Simoncelli, 2006; Sato et al., 2007).

In summary, the data suggests that as the brain programs motor commands, it also predicts the sensory consequences. Once the sensory system reports its measurements, the brain combines what it had predicted with the measurements to form a 'belief' that represents its estimate of the state of the world. This belief is then used to issue feedback responses to current actions and to adapt our internal models of the world which will guide future actions. Thus, our actions are not simply based on our current sensory observations. Rather, our actions are often based on a statistically-optimal integration of sensory observations with our predictions.

Motor adaptation: learning from sensory prediction errors

Combining predictions with observations is only useful if the predictions are generally accurate. If trial after trial there are persistent differences between predictions and observations, that is, the brain's predictions are consistently biased, then there is something wrong in these predictions. The problem of forming unbiased predictions of sensory observations is a fundamental problem of learning. When you have formed an accurate representation, i.e., a forward model of how

motor commands affect the motion of your arm or your eyes, you can apply motor commands to this internal model and (on average) accurately predict the motion that will result. However, during development, bones grow and muscle mass increases, changing the relationship between motor commands and motion of the limb. Disease can affect the strength of muscles that act on the eyes. In addition to such gradual variations, the arm's dynamics change over a shorter timescale when we grasp objects and perform manipulation. It follows that in order to maintain a desired level of performance, our brain needs to be 'robust' to these changes. This robustness may be achieved through an updating, or adaptation, of an internal model that predicts the sensory consequences of motor commands.

In the case of arm movements, there are two well studied versions of the adaptation paradigm. In one version, called *visuomotor adaptation*, the investigator introduces a perturbation that distorts the visual consequences of the motor commands but leaves the proprioceptive consequences unchanged. This is typically done by wearing prism goggles, or having people move a cursor on the screen in which the relationship between cursor position and hand position is manipulated (Krakauer et al., 1999;Krakauer et al., 2000). In another more recent version of the adaptation paradigm, called *force-field adaptation*, the investigator introduces a physical perturbation that alters both the visual and proprioceptive consequences of motor commands. This is typically done by having the volunteer hold the handle of an actuated manipulandum (a robotic arm) that can produce force on the hand that varies with hand motion (Shadmehr and Mussa-Ivaldi, 1994). This type of adaptation can also be studied by having people hold a passive manipulandum for which the weight can be adjusted (Krakauer et al., 1999), reach in a rotating room (the rotation imposes coriolis forces on the hand (Lackner and Dizio, 1994)), or even in micro-gravity in which the usual forces are removed (Lackner and DiZio, 1996).

In the case of eye movements, there is also a well studied version of the adaptation experiment. In the experiment, a target is shown and as soon as the eyes begin moving toward it, the target is extinguished and a new target is displayed (McLaughlin, 1967). As a result, the saccade completes with an endpoint error, i.e., the motor command produce the expected proprioceptive feedback from the eye muscles but an unexpected visual feedback from the retina. The current data suggests that in both the arm adaptation and saccade adaptation experiments, learning depends on the sensory prediction errors.

The oldest record of visuomotor adaptation experiment is an 1867 report by Hermann von Helmholtz. In that work, he asked subjects to point with their finger at targets while wearing prism lenses that displaced the visual field laterally. When the displacement was to the left, subjects initially had errors (an overshoot) in that direction and after some practice, they learned to compensate for the visual displacement. Helmholtz observed that as soon as the prisms were removed, subjects made erroneous movements to the right of the target. This is known as an *after-effect* of adaptation.

Nearly a century later, Held and Freedman (Held and Freedman, 1963) repeated Helmholtz's experiment with a new twist. They compared the performance of subjects when they actively moved their arm while viewing their finger through prism glasses, versus when they viewed their fingers but their arms were passively moved for them. In both cases, the subjects viewed finger motion through a prism that induced a displacement in the visual feedback. After this viewing, subjects were tested in a pointing task. Held and colleagues (HELD and Gottlieb, 1958;Held and Freedman, 1963) found that in the test session, subjects only showed after-effects if they had actively moved their hands while viewing them. In their words: "Although the passive-movement condition provided the eye with the same optical information that the active-movement condition did, the crucial connection between motor output and visual feedback was lacking". In our terminology, sensory prediction error was missing in the passive condition, as the subjects did not actively generate a movement, and therefore could not predict the sensory consequences.

There is a more recent example of visuomotor adaptation that provides striking evidence for the crucial role of sensory prediction errors. Mazzoni and Krakauer (Mazzoni and Krakauer, 2006) had people move their wrist so that the position of the index finger was coupled with the position of a cursor on a screen. There were always 8 targets on display, spanning 360°. On a given trial, one of the targets would light up and the subject would move the cursor in an out-and-back trajectory, hitting the target and then returning to the center. After a baseline familiarization period (40 trials), the experimenters imposed a 45° counter-clockwise rotation on the relationship between the cursor and finger position (early adaptation, Fig. 3A). Let us label this perturbation with r. Now, a motor command u that moved the hand in direction θ did not produce a cursor motion in the same direction, but in direction $\theta + r$. If we label the predicted sensory consequences $\hat{y} = \theta$ and the observed consequences $y = \theta + r$, then there is a sensory prediction error $y - \hat{y}$. The objective is to use this prediction error to update an estimate for $\hat{r}$.

With that estimate, for a target at direction θ^*, we can generate a motor command $u = \theta^* - \hat{r}$ to bring the cursor to the target. Indeed, after about 80 trials, in response to target at θ^* people would move their hands to $\theta^* - 40$ so the cursor would land within 5 degrees of the target (as shown in 'adaptation' subplot of Fig. 3B).

Now, Mazzoni and Krakauer (2006) took another group of naïve subjects and after they had experienced a couple of rotation trials, they simply told them: "Look, you made two movements that had large errors because we imposed a rotation that pushed you 45° counter clockwise. You can counter the error by aiming for the neighboring clockwise target." That is, simply issue the motor command $u = \theta^* - 45$ and as a consequence, the cursor will move at direction θ and land at the target. Indeed, the subjects followed this strategy: on the very next trial, all the error dropped to zero (strategy group, Fig. 3B). However, something very interesting happened: as the trials continued, the errors gradually grew! What's more, the rate of change in the errors in this 'strategy' group was exactly the same as the rate of change in the regular adaptation paradigm.

To explain this, Mazzoni and Krakauer (2006) hypothesized that on trial 43, when the subjects in the strategy group were producing the motor commands that brought the cursor to the target, there was still a discrepancy between the predicted and observed sensory consequences of motor commands $y - \hat{y}$. This is because whereas *explicitly* they had been told of the perturbation, *implicitly* their estimate was still around zero, $\hat{r} \approx 0$. The "learning curve" (over which the performance errors grew in the strategy group) tells us how quickly the implicit estimates became accurate. Interestingly, this occurred gradually over the course of about 50 trials, which closely matches the time course of normal adaptation (compare panels A & B in Fig. 3), suggesting that the ability to predict the consequences of our actions and the ability to control our actions improve at the same rate. If accurate prediction preceded accurate control, the learning curve over which performance errors increased in the strategy group would have been more rapid than the time course of normal adaptation. This finding challenges the notion that accurate prediction (i.e., learning of a forward model) precedes the ability to control ones actions (Flanagan et al., 2003), and suggests instead that accurate prediction may be the limiting factor for accurate control.

However, there is another way to view the process of motor adaptation. As we perform a movement that experiences an error, we sometimes have the opportunity to correct it (for

example, via reflexive pathways during that movement, or via corrective movements that occur later). These motor corrections may act as a teaching signal for the brain (Miles and Lisberger, 1981;Kawato, 1996). For example, when arm movements are perturbed with unexpected forces, reflex pathways respond to partially compensate for the sensory prediction errors (Thoroughman and Shadmehr 1999). When a visual target is moved during a saccade, there is a second saccade that brings the eyes to the new target position. One way to think of motor learning is to imagine that the motor commands that corrected the movement might be added with a slight time advance to the motor commands that initiate the next movement. The error feedback learning theory of the cerebellum relies on this motor correction (Kawato 1996).

To address this question, experiments have attempted to dissociate the effect of error-driven motor corrections from error signals themselves. In saccade adaptation experiments this has been done by presenting endpoint errors (sensory prediction errors) under conditions that reduce motor corrections (Wallman and Fuchs, 1998;Noto and Robinson, 2001). In reach adaptation paradigms this has been done by making a rapid movement so that there is a reduced possibility of correcting that movement (Tseng et al., 2007). All three experiments demonstrate that adaptation of the motor commands is driven directly by error signals rather than error corrections.

The multiple timescales of adaptation

Recent work has shown that the interactions between adaptive processes which learn and forget on different timescales underlie several key features of motor adaptation (Smith et al., 2006). The idea that the process of memory formation proceeds on multiple timescales is, of course, not new as multiple timescales have long been observed in learning curves and forgetting curves (Rubin and Wenzel, 1996;Scheidt et al., 2000;Wixted, 2004). However, the ability to understand the natural interactions between different adaptive processes opens a window onto understanding several seemingly complex learning phenomena.

One such phenomenon is spontaneous recovery. Spontaneous recovery of memory refers to the natural reemergence of a learned response after that response has been extinguished through extinction training. This phenomenon has been observed in a wide range of paradigms (Rescorla, 2004;Kojima et al., 2004;Stollhoff et al., 2005). Our recent work proposed a mechanism that explains spontaneous recovery as an interaction between the decay of two adaptive processes (Fig. 4). If two processes participate in learning from error, one with a faster time course than the

other, then both processes need not be extinguished for the overall learning to disappear. Instead, learning can be extinguished if these processes come to cancel each other's effects during the extinction training. If this is the case, then the extinction training would leave two opposing memory traces that compete with each other and decay at different rates. Since the extinction training is generally quite rapid, the faster process would be expected to oppose the initial learning, and its decay would reveal the incomplete extinction of the slower process – corresponding to a spontaneous recovery of the overall learning. Decay of this slower process would eventually extinguish the spontaneous recovery.

In recent studies, the pattern of spontaneous recovery was shown to correspond to the intrinsic decay rates of two adaptive processes in exactly this way for motor adaptation in both arm (Smith et al., 2006) and eye movements (Ethier et al., 2008). In the first study, the adaptation to a velocity-dependent force-field was extinguished by exposure to the opposite force-field. When motor errors were subsequently clamped to zero, the adaptive response to the original force-field spontaneously reemerged with a fast exponential time course and then died away with a slower exponential time course as would be predicted by this model. In the second study, the same pattern of spontaneous recovery was replicated for saccade adaptation.

The basic idea that extinction results when the fast component of learning cancels the effects of slower components provides a possible explanation for two other key features of learning: savings and anterograde interference. If the original response is retrained after extinction, this relearning is noticeably faster than the initial learning despite the same initial performance level and the same training procedure – a phenomenon known as savings or facilitation. However if a response opposite to the initial response is trained after extinction, this opposite learning is slower than the initial learning – a phenomenon known as anterograde interference or forward interference. Because the fast process responds more quickly to training than the slow process, subsequent training compatible with the initial learning will benefit from the bias of the slow learning process toward the initial learning. In contrast, subsequent training which opposes the initial learning will suffer from this bias.

It should be noted, however, that neither savings nor anterograde interference can be fully explained by simple interactions between adaptive processes. A recent study showed that a certain amount of savings persists even when extinction is carried out gradually under conditions in which the individual adaptive processes should be extinguished (Zarahn et al., 2008). The

mechanism underlying this component of savings is not clear, but it may result from a general improvement in the rate of learning or from a memory of the initial adaptation that is not expressed until it is recalled. In an earlier study, catastrophic anterograde interference occurred after delays of several days, suggesting the operation of contextual effects rather than persistence of adaptive processes (Krakauer et al., 2005). Another studied found that when training was followed by 'reverse-training' until extinction, the spontaneous recovery was so strong as to suggest that during reverse-training, there was effectively no unlearning, but rather only an instantiation of a new and competing fast adaptive memory (Criscimagna-Hemminger and Shadmehr, 2008).

Whereas interactions between the fast and slow processes account for initial learning, spontaneous recovery, and some forms of interference, it appears that only one of these processes provides a gateway to long-term memory formation. 24-hour retention levels assessed after various amounts of initial training (i.e. at different points of the learning curve) do not directly reflect the adaptation level achieved during the initial training period, but instead reflect the predicted level of the slow component of adaptation as shown in figure 5 (Joiner and Smith, 2008). This suggests that maximizing the long-term benefit of a training session does not necessarily come from maximizing the overall level of learning, but rather from maximizing the amount of learning achieved by a single constituent, the slow process. Thus, a simple interaction between a fast and a slow adaptive process appears to explain not only the shapes of initial learning curves, but also the phenomena of spontaneous recovery, anterograde interference, and patterns of 24-hour retention in motor adaptation.

Eye movements also display multiple time scales of adaptation. For example, eye velocity in optokinetic nystagmus (involuntary eye movements in response to continuous movement of the visual field) is characterized by two components: a rapid rise followed by a slower increase to steady state (Cohen et al., 1977). Interestingly, the rapid rise in eye velocity has been shown to be specifically affected by particular neural lesions (Zee et al., 1987), suggesting that these time scales may have distinct neuroanatomical bases. This idea is supported by data indicating that during saccade adaptation in patients with cerebellar cortical damage, there is a profound loss in the fast timescales of adaptation but less impairment in the slower adaptive processes (Xu-Wilson et al., 2009a). In reaching movements, transcranial magnetic stimulation of the lateral posterior parietal cortex (Della-Maggiore et al., 2004) or the motor cortex (Richardson et al., 2006;Hadipour-Niktarash et al., 2007) appears to specifically affect the slower adaptive

processes. Interestingly, stimulation of posterior parietal cortex appears to have a stronger effect than stimulation of motor cortex. Imaging studies also suggest that there may be distinct neural networks associated with the fast and slow adaptive processes (Krakauer et al., 2004;Tunik et al., 2007)

Limitations

The framework that we focused on in this review is one in which the brain learns to accurately predict the sensory consequences of motor commands. Three important and pertinent problems in motor learning have not been discussed here: how to discern common structure in the dynamics of the tasks in which one is engaged in (e.g., uncovering structure) how to produce motor commands that maximize some measure of performance (e.g., maximizing reward while minimizing effort, as in optimal control), and how the brain decides when to form new memories and when to modify existing ones.

The first problem is one of discovering the structure of the system that one is controlling. For example, inertial objects share a common structure in which acceleration and velocity are linearly separable quantities in the function that relates motion to forces. An intelligent system that interacts with a large number of inertial objects would benefit from discovering this structure, as it would allow it to learn control by simply adjusting a small number of parameters, rather than searching in the high-dimensional space of all possible motor commands (Hwang et al., 2006;Braun et al., 2009). It is possible that in the short timescale of training in a laboratory, subjects use internal models with structures that are similar to that of their body and merely 'tweak' the parameters of these models to adapt to the imposed perturbations (Kording et al., 2007). Another way to say this is that the mechanisms involved in correcting movements in response to self-generated perturbations (natural variability in the motor commands that move our body parts) are the same as those co-opted when people adapt to externally imposed perturbations. However, long-term learning likely involves building a structural model specific to the task at hand, a process which would allow specialization and expertise that goes beyond parameter estimation with a generic model (Reis et al., 2009).

The second problem is that even if one has a perfect model that predicts the sensory consequences of motor commands, one still has to find the motor commands that maximize performance. This is the general problem of optimal control, i.e., given a relationship between motor commands and

sensory consequences, find a feedback control law that achieves the greatest amount of reward (or some other performance measure) at the least amount of effort (Todorov and Jordan, 2002). Evidence for this process comes from experiments that find during adaptation to a perturbation, motor commands do not return performance to a baseline trajectory, but rather a trajectory that maximizes performance while minimizing a measure of effort (Uno et al., 1989;Emken et al., 2007;Izawa et al., 2008).

The third problem is that if adaptive changes to the mechanisms for predicting the consequences of our actions and for controlling our actions occur continuously as described here, it is difficult to imagine that each movement leads to a new memory that can be recalled. On the other hand, we know that multiple discrete memories can be formed and consolidated. Understanding the nature of the contexts that lead to the formation and recall of motor memories is clearly important (Gandolfo et al., 1996;Krakauer et al., 2006;Krakauer, 2009;Cothros et al., 2009). State, as a physicist would define it – position and velocity of motion - has been shown to be an important context for the formation and recall of motor memories, but multiple memories can be associated with the same motion states. What other contexts are important, and how do some contexts come to be more important than others?

Summary and conclusions

The motor system needs to be able to adjust for the both the presence of noise and delay in sensory feedback, and for changes in the body and the world that alter the relationship between motor commands and their sensory consequences. The common solution to these two problems is a forward model. Forward models, possibly located in the cerebellum, receive a copy of the outgoing motor command and generate a prediction of the expected sensory consequence at very short latency. This output in sensory coordinates can be used to make fast trajectory corrections with the pre-existing controller before true sensory feedback is available, and can be integrated with true sensory feedback to optimize state estimates and enhance perception. Forward models are only useful if they are accurate and evidence suggests that accuracy is maintained through adaptive processes driven by sensory prediction errors.
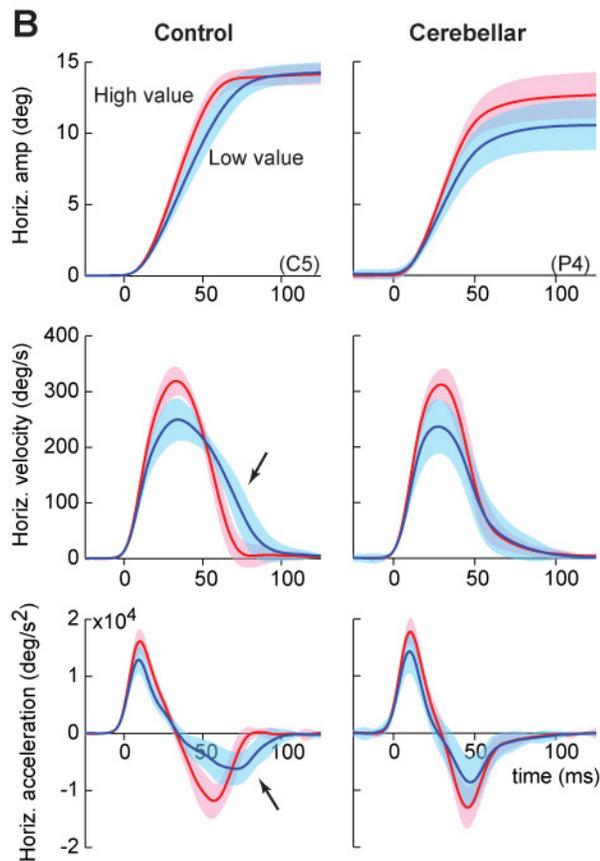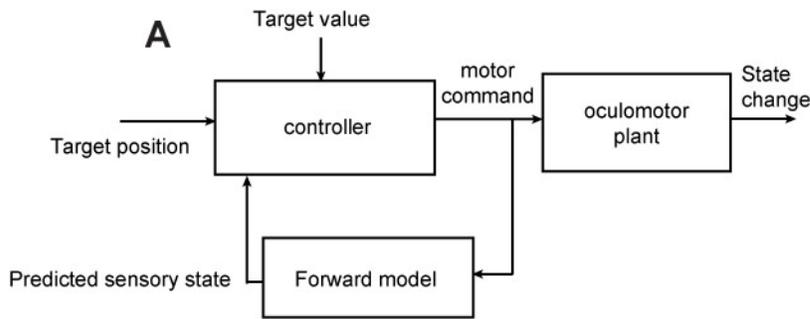
Figure 1. Control of movements without sensory feedback is difficult because there is variability in the motor commands. The brain appears to maintain accuracy by using a forward model that predicts the sensory consequences of motor commands. **A)** A highly stylized view of the process of generating a saccade. The motor commands depend not just on the position of the saccade target, but also the internal value associated with that movement. For example, saccades tend to be slower toward stimuli that have a lower value. **B)** Examples of saccades from a healthy subject and a patient with degeneration of the cerebellum. With repeated presentation of a visual target, the saccade target is devalued, and the motor commands that initiate the saccade become smaller, generating slower velocities and accelerations in both the healthy subject and the cerebellar patient (blue lines). In response to this variability, the healthy brain produces motor commands late in the saccade to maintain accuracy, bringing the eyes to the target. In contrast, the patient with cerebellar damage cannot correct for the reductions in the motor commands that

20

initiated the saccade. As a result, the saccades of the cerebellar patient fall short of the target. The figure in part B is reproduced from (Xu-Wilson et al., 2009a).
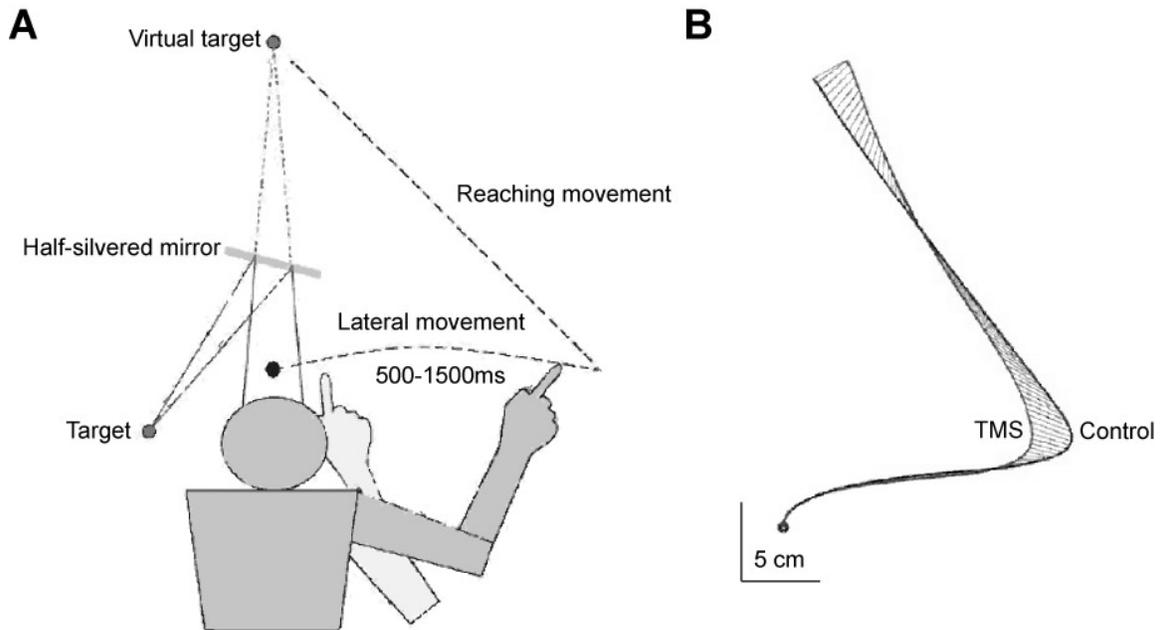

Figure 2. Control of movements with sensory feedback is difficult because there feedback is time delayed. The brain appears to maintain accuracy by using a forward model that predicts the sensory consequences of motor commands. A) Volunteers were asked to move their hand laterally until they heard a tone, at which point they would reach toward a target. The cerebellum was disrupted via a TMS pulse soon after the tone. B) Reach trajectories from a movement in which TMS was applied to the cerebellum, and a movement in which no TMS was applied. Application of TMS produces a movement in which the estimate of the state of the arm appears to be delayed with respect to its actual state, resulting in missing the target to the right. From (Miall et al., 2007).
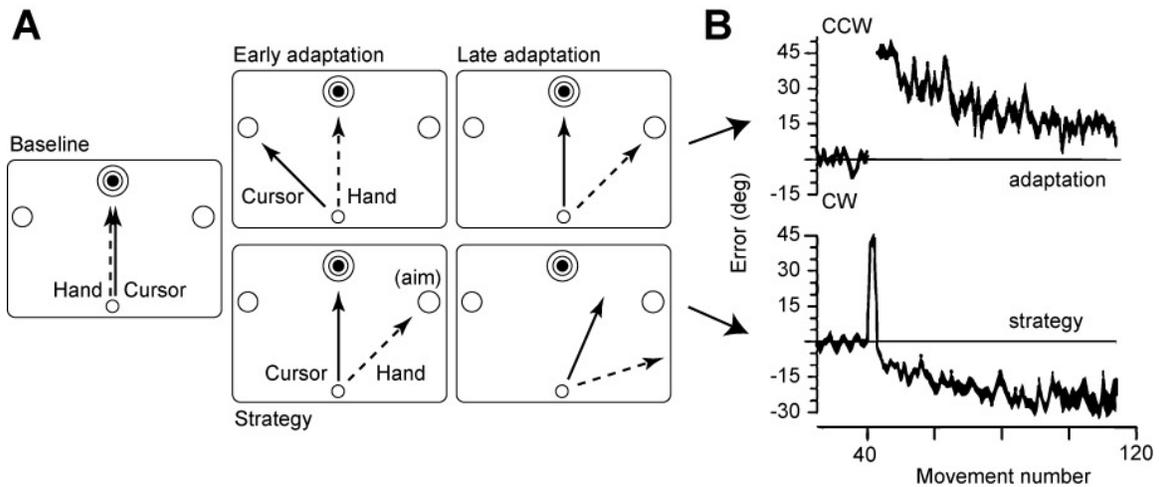
Figure 3. An example of learning from sensory prediction errors during visuomotor adaptation. **A**) Subjects were asked to make an out-and-back motion with their hand so a cursor was moved to one of 8 targets. In the baseline condition, hand motion and cursor motion were congruent. In the adaptation condition, a 45^o rotation was imposed on the motion of the cursor and the hand. In the adaptation group (top two plots), the subjects gradually learned to move their hand in a way that compensated for the rotation. In the strategy group (bottom two plots), after two movements subjects were told about the perturbation and asked to simply aim to the neighboring target. **B**) Endpoint errors in the adaptation and strategy groups. The strategy group immediately compensated for the endpoint errors, but paradoxically, the errors gradually grew. The rate of change of errors in the strategy and adaptation groups was the same. The rapid initial improvement is due to learning in the explicit memory system, whereas the gradual learning that follows is due to an implicit system. From (Mazzoni and Krakauer, 2006).
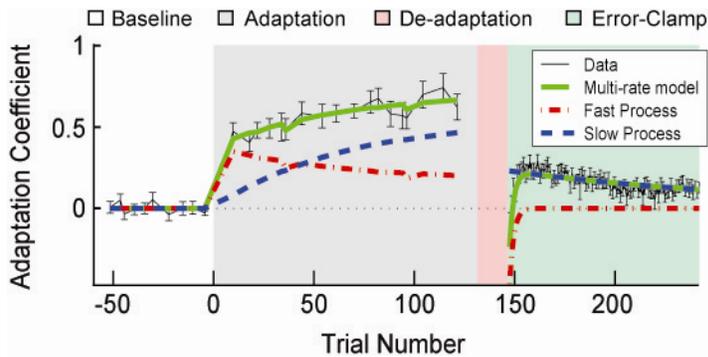
Figure 4. Spontaneous recovery in motor adaptation is explained by a two-process multi-rate model. After a baseline period (no shading), participants were trained to adapt to a velocity-dependent force-field (gray shading), de-adapted by exposure to the opposite force-field (red shading), and tested for recovery under conditions in which lateral errors were held to zero. Adaptation levels during the recovery period were measured on error-clamp trials during which lateral errors were clamped at zero to prevent error-driven learning and so that the feedforward patterns of lateral forces could be directly measured without contamination from error-driven feedback corrections or inertial interactions. These error-clamp trials were also occasionally interspersed into the baseline and training periods to measure the initial learning curve. The multi-rate model reproduces the double exponential pattern in the learning curve, and not only the amount, but the shape of the spontaneous recovery profile. According to the model, this shape emerges because of the superposition of a rapid up-going decay of the fast learning process and a prolonged down-going decay of the slow learning process. From (Smith et al., 2006).
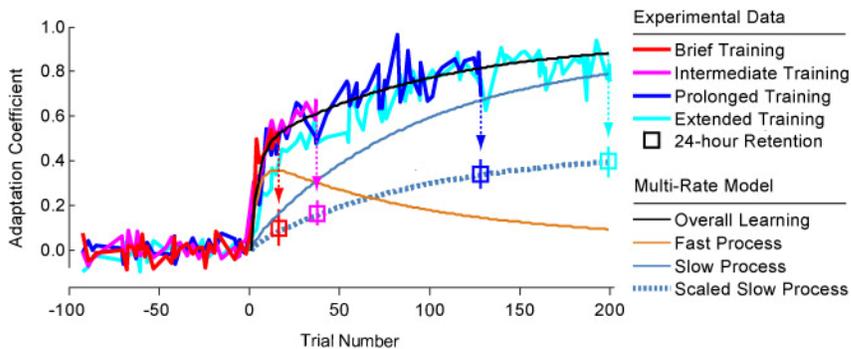
Figure 5. Long-term retention in motor adaptation is explained by the level of a slow component of motor adaptation. Four groups of participants were tested for 24-hour retention after different amounts of training with a velocity-dependent force-field. The learning curves during the training period itself are plotted with lines, and the 24-hour retention levels are plotted as squares (mean ± SEM). The pattern of 24-hour retention has essentially the same shape as the slow learning process (r=0.99), suggesting that this process serves as a gateway to long-term retention. From (Joiner and Smith, 2008).

Reference List

Ariff GD, Donchin O, Nanayakkara T, Shadmehr R (2002) A Real-time state predictor in motor control: Study of saccadic eye movements during unseen reaching movements. J Neurosci 22:7721-7729.

Bhushan N, Shadmehr R (1999) Computational architecture of human adaptive control during learning of reaching movements in force fields. Biol Cybern 81:39-60.

Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Motor task variation induces structural learning. Curr Biol 19:352-357.

Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. J Neurosci 28:2804-2813.

Cohen B, Matsuo V, Raphan T (1977) Quantitative analysis of the velocity characteristics of optokinetic nystagmus and optokinetic after-nystagmus. J Physiol 270:321-344.

Cothros N, Wong J, Gribble PL (2009) Visual cues signaling object grasp reduce interference in motor learning. J Neurophysiol. in press.

Criscimagna-Hemminger SE, Shadmehr R (2008) Consolidation patterns of human motor memory. J Neurosci 28:9610-9618.

Della-Maggiore V, Malfait N, Ostry DJ, Paus T (2004) Stimulation of the posterior parietal cortex interferes with arm trajectory adjustments during the learning of new dynamics. J Neurosci 24:9971-9976.

Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ (2007) Motor adaptation as a greedy optimization of error and effort. J Neurophysiol 97:3997-4006.

Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. J Neurophysiol 99:2577-2583.

Flanagan JR, Vetter P, Johansson RS, Wolpert DM (2003) Prediction precedes control in motor learning. Curr Biol 13:146-150.

Fuchs AF, Binder MD (1983) Fatigue resistance of human extraocular muscles. J Neurophysiol 49:28-34.

Gandolfo F, Mussa-Ivaldi FA, Bizzi E (1996) Motor learning by field approximation. Proc Natl Acad Sci USA 93:3843-3846.

Golla H, Tziridis K, Haarmeier T, Catz N, Barash S, Thier P (2008) Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. Eur J Neurosci 27:132-144.

Guthrie BL, Porter JD, Sparks DL (1983) Corollary discharge provides accurate eye position information to the oculomotor system. Science 221:1193-1195.

Haarmeier T, Thier P, Repnow M, Petersen D (1997) False perception of motion in a patient who cannot compensate for eye movements. Nature 389:849-852.

Hadipour-Niktarash A, Lee CK, Desmond JE, Shadmehr R (2007) Impairment of retention but not acquisition of a visuomotor skill through time-dependent disruption of primary motor cortex. J Neurosci 27:13413-13419.

Held R, Freedman SJ (1963) Plasticity in human sensorimotor control. Science 142:455-462.

HELD R, Gottlieb N (1958) Technique for studying adaptation to disarranged hand-eye coordination. Perceptual Motor Skills 8:83-86.

Hwang EJ, Smith MA, Shadmehr R (2006) Adaptation and generalization in acceleration-dependent force fields. Exp Brain Res 169:496-506.

Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. J Neurosci 28:2883-2891.

Izawa J, Shadmehr R (2008) Online processing of uncertain information in visuomotor control. J Neurosci 28:11360-11368.

Joiner WM, Smith MA (2008) Long-term retention explained by a model of short-term learning in the adaptive control of reaching. J Neurophysiol 100:2948-2955.

Kalman RE (1960) A new approach to linear filtering and prediction problems. Trans ASME J Basic Engineering 82 (series D):35-45.

Kawato M (1996) Learning internal models of the motor apparatus. In: The acquisition of motor behavior in vertebrates (Bloedel JR, Ebner TJ, Wise SP, eds), pp 409-430. Cambridge, MA: MIT Press.

Keller EL, Robinson DA (1971) Absence of a stretch reflex in extraocular muscles of the monkey. J Neurophysiol 34:908-919.

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. J Neurosci 24:7531-7539.

Kording KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. J Neurophysiol 92:3161-3165.

Kording KP, Tenenbaum JB, Shadmehr R (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat Neurosci 10:779-786.

Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Krakauer JW (2009) Motor learning and consolidation: the case of visuomotor rotation. Adv Exp Med Biol 629:405-421.

Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. J Neurosci 25:473-478.

Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. Nature Neurosci 2:1026-1031.

Krakauer JW, Ghilardi MF, Mentis M, Barnes A, Veytsman M, Eidelberg D, Ghez C (2004) Differential cortical and subcortical activations in learning rotations and gains for reaching: a PET study. J Neurophysiol 91:924-933.

Krakauer JW, Mazzoni P, Ghazizadeh A, Ravindran R, Shadmehr R (2006) Generalization of Motor Learning Depends on the History of Prior Action. PLoS Biol 4:e316.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. J Neurosci 20:8916-8924.

Kurtzer IL, Pruszynski JA, Scott SH (2008) Long-latency reflexes of the human arm reflect an internal model of limb dynamics. Curr Biol 18:449-453.

Lackner JR, Dizio P (1994) Rapid adaptation to coriolis force perturbations of arm trajectory. J Neurophysiol 72:299-313.

Lackner JR, DiZio P (1996) Motor function in microgravity: movement in weightlessness. Curr Opin Neurobiol 6:744-750.

Lacquaniti F, Soechting JF (1984) Behavior of the stretch reflex in a multi-jointed limb. Brain Res 311:161-166.

Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. J Neurosci 26:3642-3645.

McLaughlin S (1967) Parametric adjustment in saccadic eye movements. Percept Psychophys 2:359-362.

Mehta B, Schaal S (2002) Forward models in visuomotor control. J Neurophysiol 88:942-953.

Miall RC, Christensen LOD, Owen C, Stanley J (2007) Disruption of state estimation in the human lateral cerebellum. PLoS Biol 5:e316.

Miall RC, Wolpert DM (1996) Forward models for physiological motor control. Neural Networks 9:1265-1279.

Miles FA, Lisberger SG (1981) The "error" signals subserving adaptive gain control in the primate vestibulo-ocular reflex. Ann N Y Acad Sci 374:513-525.

Noto CT, Robinson FR (2001) Visual error is the stimulus for saccade gain adaptation. Brain Res Cogn Brain Res 12:301-305.

Pasalar S, Roitman AV, Durfee WK, Ebner TJ (2006) Force field effects on cerebellar Purkinje cell discharge with implications for internal models. Nat Neurosci.

Reis J, Schambra HM, Cohen LG, Buch ER, Fritsch B, Zarahn E, Celnik PA, Krakauer JW (2009) Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. Proc Natl Acad Sci U S A 106:1590-1595.

Rescorla RA (2004) Spontaneous recovery. Learn Mem 11:501-509.

Richardson AG, Overduin SA, Valero-Cabre A, Padoa-Schioppa C, Pascual-Leone A, Bizzi E, Press DZ (2006) Disruption of primary motor cortex before learning impairs memory of movement dynamics. J Neurosci 26:12466-12470.

Robinson DA (1975) Oculomotor control signals. In: Basic Mechanisms of Ocular Motility and Their Clinical Implications (BachyRita P, Lennerstrand G, eds), pp 337-374. Oxford, UK: Pergamon.

Rottach KG, Das VE, Wohlgemuth W, Zivotofsky AZ, Leigh RJ (1998) Properties of horizontal saccades accompanied by blinks. J Neurophysiol 79:2895-2902.

Rubin DC, Wenzel AE (1996) One hundred years of forgetting: a quantitative description of retention. Psychol Rev 103:734-760.

Saunders JA, Knill DC (2003) Humans use continuous visual feedback from the hand to control fast reaching movements. Exp Brain Res 152:341-352.

Saunders JA, Knill DC (2004) Visual feedback control of hand movements. J Neurosci 24:3223-3234.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. J Neurophysiol 84:853-862.

Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. Exp Brain Res 185:359-381.

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14:3208-3224.

Smit AC, Van Gisbergen JA, Cools AR (1987) A parametric analysis of human saccades in different experimental paradigms. Vision Res 27:1745-1762.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179.

Snyder LH, Calton JL, Dickinson AR, Lawrence BM (2002) Eye-hand coordination: saccades are faster when accompanied by a coordinated arm movement. J Neurophysiol 87:2279-2286.

Soechting JF, Lacquaniti F (1988) Quantitative evaluation of the electromyographic responses to multidirectional load perturbations of the human arm. J Neurophysiol 59:1296-1313.

Sommer MA, Wurtz RH (2002) A pathway in primate brain for internal monitoring of movements. Science 296:1480-1482.

Sommer MA, Wurtz RH (2006) Influence of the thalamus on spatial visual processing in frontal cortex. Nature 444:374-377.

Stollhoff N, Menzel R, Eisenhardt D (2005) Spontaneous recovery from extinction depends on the reconsolidation of the acquisition memory in an appetitive learning paradigm in the honeybee (Apis mellifera). J Neurosci 25:4485-4492.

Straube A, Fuchs AF, Usher S, Robinson FR (1997) Characteristics of saccadic gain adaptation in rhesus macaques. J Neurophysiol 77:874-895.

Strick PL (1978) Cerebellar involvement in volitional muscle responses to load change. In: Cerebral motor control in man: long loop mechanisms (Desmedt JE, ed), pp 85-93. Basel, Switzerland: Karger.

Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. Exp Brain Res 142:284-291.

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226-1235.

Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. J Neurophysiol 98:54-62.

Tunik E, Schmitt PJ, Grafton ST (2007) BOLD coherence reveals segregated functional neural interactions when adapting to distinct torque perturbations. J Neurophysiol 97:2107-2120.

Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. Biol Cybern 61:89-101.

van Donkelaar P (1997) Eye-hand interactions during goal-directed pointing movements. Neuroreport 8:2139-2142.

Vaziri S, Diedrichsen J, Shadmehr R (2006) Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. J Neurosci 26:4188-4197.

Wagner MJ, Smith MA (2008) Shared internal models for feedforward and feedback control. J Neurosci 28:10663-10673.

Wallman J, Fuchs AF (1998) Saccadic gain modification: visual error drives motor adaptation. J Neurophysiol 80:2405-2416.

Wixted JT (2004) On Common Ground: Jost's (1897) law of forgetting and Ribot's (1881) law of retrograde amnesia. Psychol Rev 111:864-879.

Xu-Wilson M, Chen-Harris H, Zee DS, Shadmehr R (2009a) Cerebellar contributions to adaptive control of saccades in humans. J Neurosci in press.

Xu-Wilson M, Zee DS, Shadmehr R (2009b) The intrinsic value of visual information affects saccade velocities. Exp Brain Res 196:475-481.

Zarahn E, Weston GD, Liang J, Mazzoni P, Krakauer JW (2008) Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. J Neurophysiol 100:2537-2548.

Zee DS, Tusa RJ, Herdman SJ, Butler PH, Gucer G (1987) Effects of occipital lobectomy upon eye movements in primate. J Neurophysiol 58:883-907.